global and local matching of the rendered motion vector field to the corresponding motion vector field extracted from the video stream via a (standard or customized) motion estimation algorithm, thereby enabling directed adjustment of model parameters;

a hierarchical sequencing of model parameter class refinement when applying said method; and

the application of reverse mapping to attain approximate values for un refined model parameter classes.

ABSTRACT

This application concerns an enhancement of the render, match, and refine (RMR) method [0002] for 3D scene model refinement. A method is provided that allows for the directed refinement of camera and geometry associated parameters of the 3D scene model in a manner that is independent of lighting, color, and texture. The method is based on the matching of a rendered motion vector field to a motion vector field estimated from a recorded video stream.

BACKGROUND OF THE INVENTION

[0001] Automated 3D scene model refinement based on camera recordings has at least three application domains: computer vision, video compression, and 3D scene reconstruction.

[0002] The render, match, and refine (RMR) method for 3D scene model refinement involves rendering a 3D model to a 2D frame buffer, or a series of 2D frames, and comparing these to images or video streams recorded using one or more cameras. The mismatch between the rendered and recorded frames is subsequently used to direct the refinement of the scene model. The intended result is that on iterative application of this procedure, the 3D scene model elements (viewpoint, vertices, NURBS, lighting, textures, etc.) will converge on an optimal description of the recorded actual scene. The field of analogous model-based methods of which the RMR method is part is known as CAD-based vision.

[0003] Many implementations of 3D to 2D rendering pipelines exist. These perform the various steps involved in calculating 2D frames from a 3D scene model. When motion is modeled,